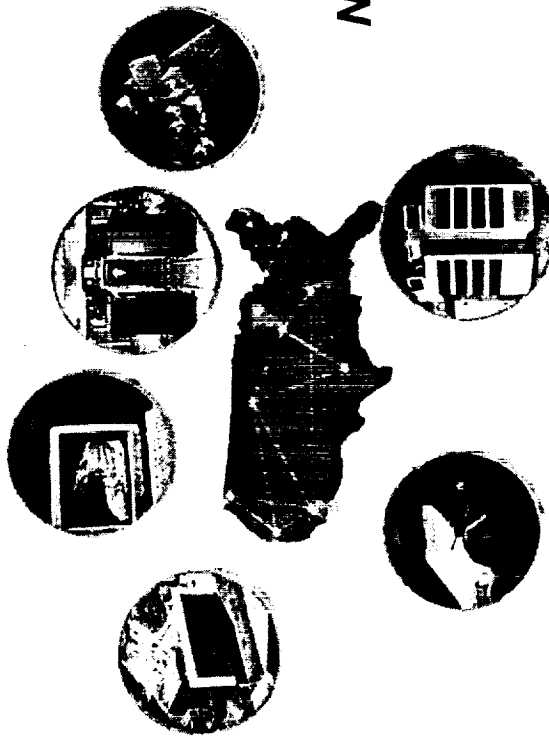




NASA's Information Power Grid

Large Scale Distributed Computing and Data Management



William E. Johnston, Project Manager

Arsi Vaziri, Deputy Project Manager

Tom Hinke, Deployment Task Manager

Leigh Ann Tanner, Implementation Manager

*NASA Advanced Supercomputing (NAS) Division
NASA Ames Research Center*

William J. Feiereisen, Division Chief

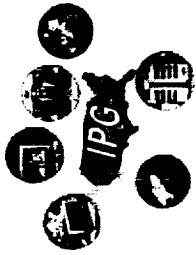
William Thigpen, Engineering Branch Chief,

<http://www.ipg.nasa.gov>



Structure of the Talk

- **Motivation**
- **What are Grids?**
- **What is Being Done with Grids?**
- **Approach to NASA's Information Power Grid**
- **State of IPG**
- **Lessons Learned for Building Large-Scale Grids**
- **Directory Services Technology Issues for Large-Scale Grids**

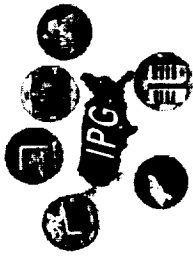




➤ ***Motivation for IPG***

¥ Large-scale science and engineering are done through the interaction of people, heterogeneous computing resources, information systems, and instruments, all of which are geographically and organizationally dispersed.

¥ *The overall motivation for Grids is to facilitate the routine interactions of these resources in order to support large-scale science and engineering.*

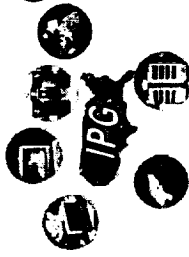




Applications Motivating Grid Computing Environments



- ¥ Multi-disciplinary simulations provide a good example of a class of applications that are very likely to require aggregation of widely distributed computing, data, and intellectual resources.
- ¥ Such simulations —e.g. whole system aircraft simulation and whole system living cell simulation —*require integrating applications and data that are developed by different teams of researchers frequently in different locations.*
- ¥ *The research teams are the only ones that have the expertise to maintain and improve the simulation code and/or the body of experimental data that drives the simulations.* This results in an inherently distributed computing and data management environment.



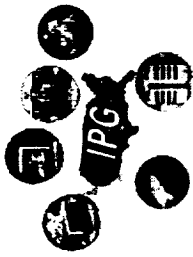


Applications Motivating Grid Computing Environments

¥ Consider a vision for Aviation Safety: How do we simulate the entire commercial airspace of the country?

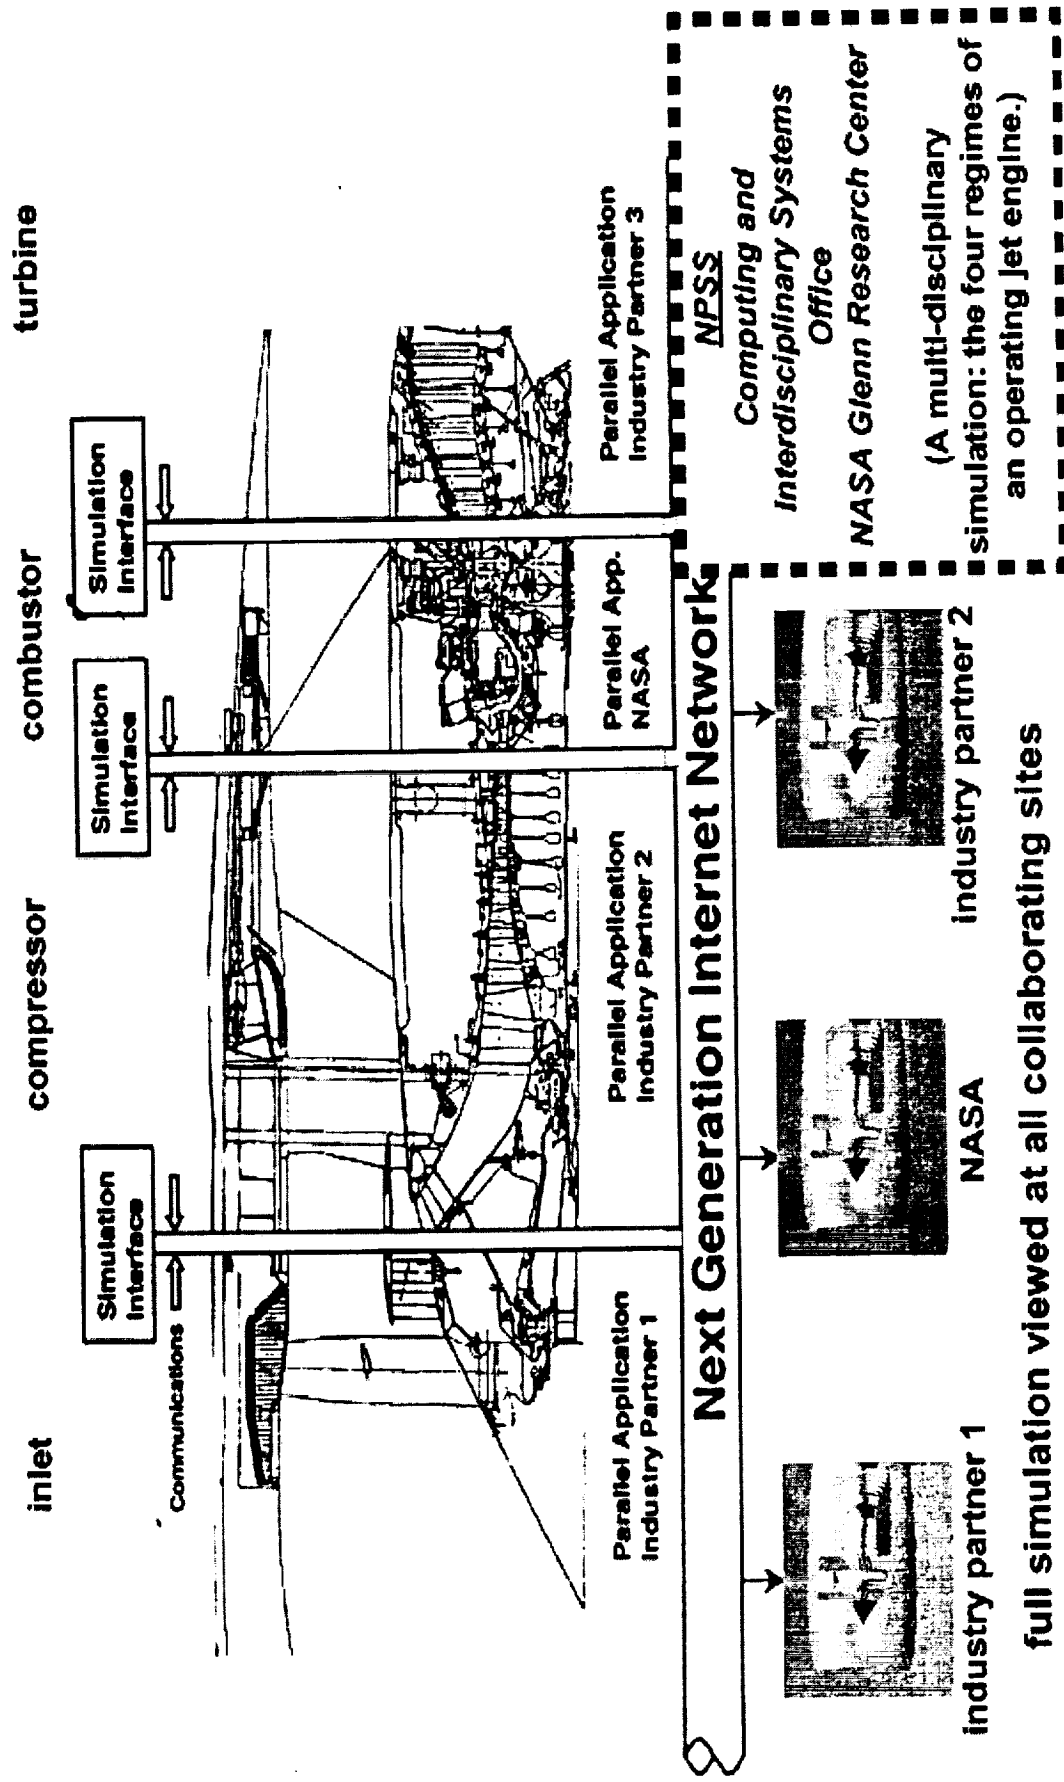
—Yuri Gawdiak and Bill McDermott, NASA Ames,
John Lytle and Gregory Follen, NASA Glenn

¥ This vision is being approached through a set of increasingly complex and computationally intensive distributed system integrations





Multi-disciplinary Simulations: Aviation Safety

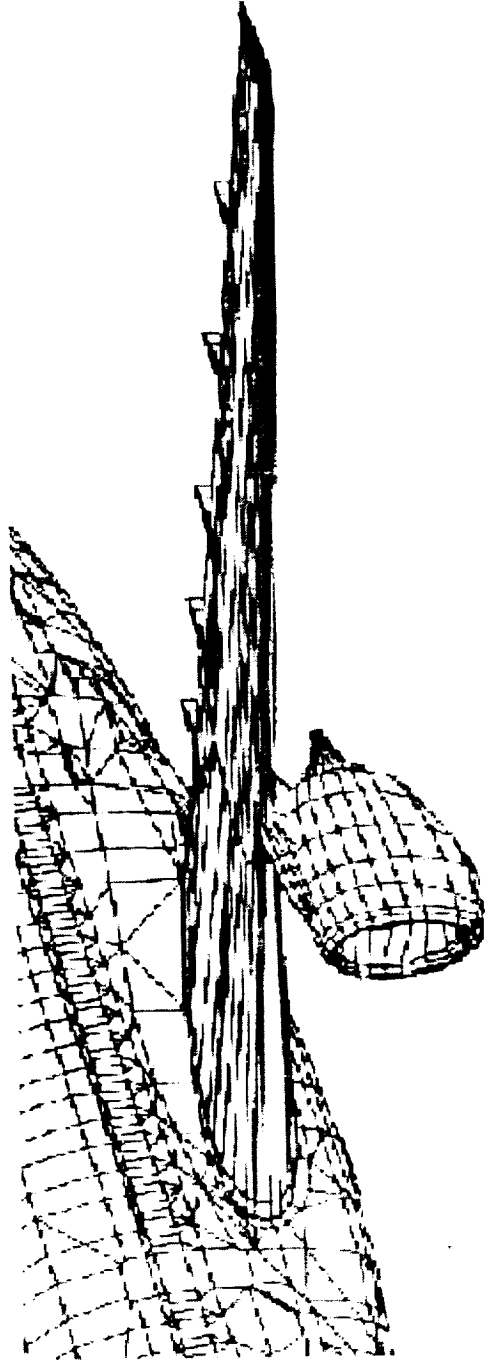


full simulation viewed at all collaborating sites

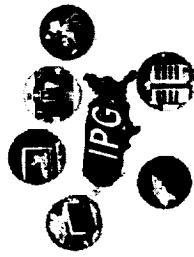
*Component simulations are combined to get
sub-system simulations*



Multi-disciplinary Simulations: Aviation Safety



Multiple sub-systems, e.g. a wing lift model operating at NASA Ames and a turbo-machine model operating at NASA Glenn, are combined using an application framework like NPSS to manage the multiple models and to use IPG services to coordinate computing and data storage systems across NASA.





Multi-disciplinary Simulations: Aviation Safety



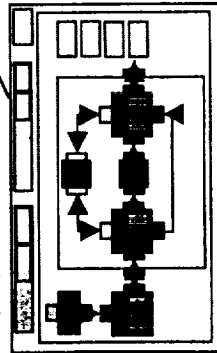
Wing Models

✱Lift Capabilities
✱Drag Capabilities
✱Responsiveness

Airframe Models

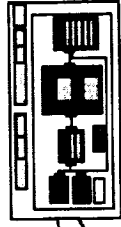


Crew Capabilities
- accuracy
- perception
- stamina
- re-action times



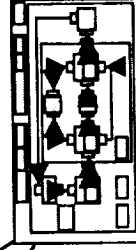
Human Models

Stabilizer Models



✱Deflection capabilities
✱Responsiveness

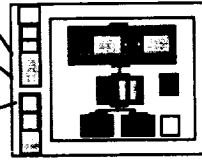
Engine Models



✱Thrust performance
✱Reverse Thrust performance
✱Responsiveness
✱Fuel Consumption

✱Braking performance
✱Steering capabilities
✱Traction
✱Dampening capabilities

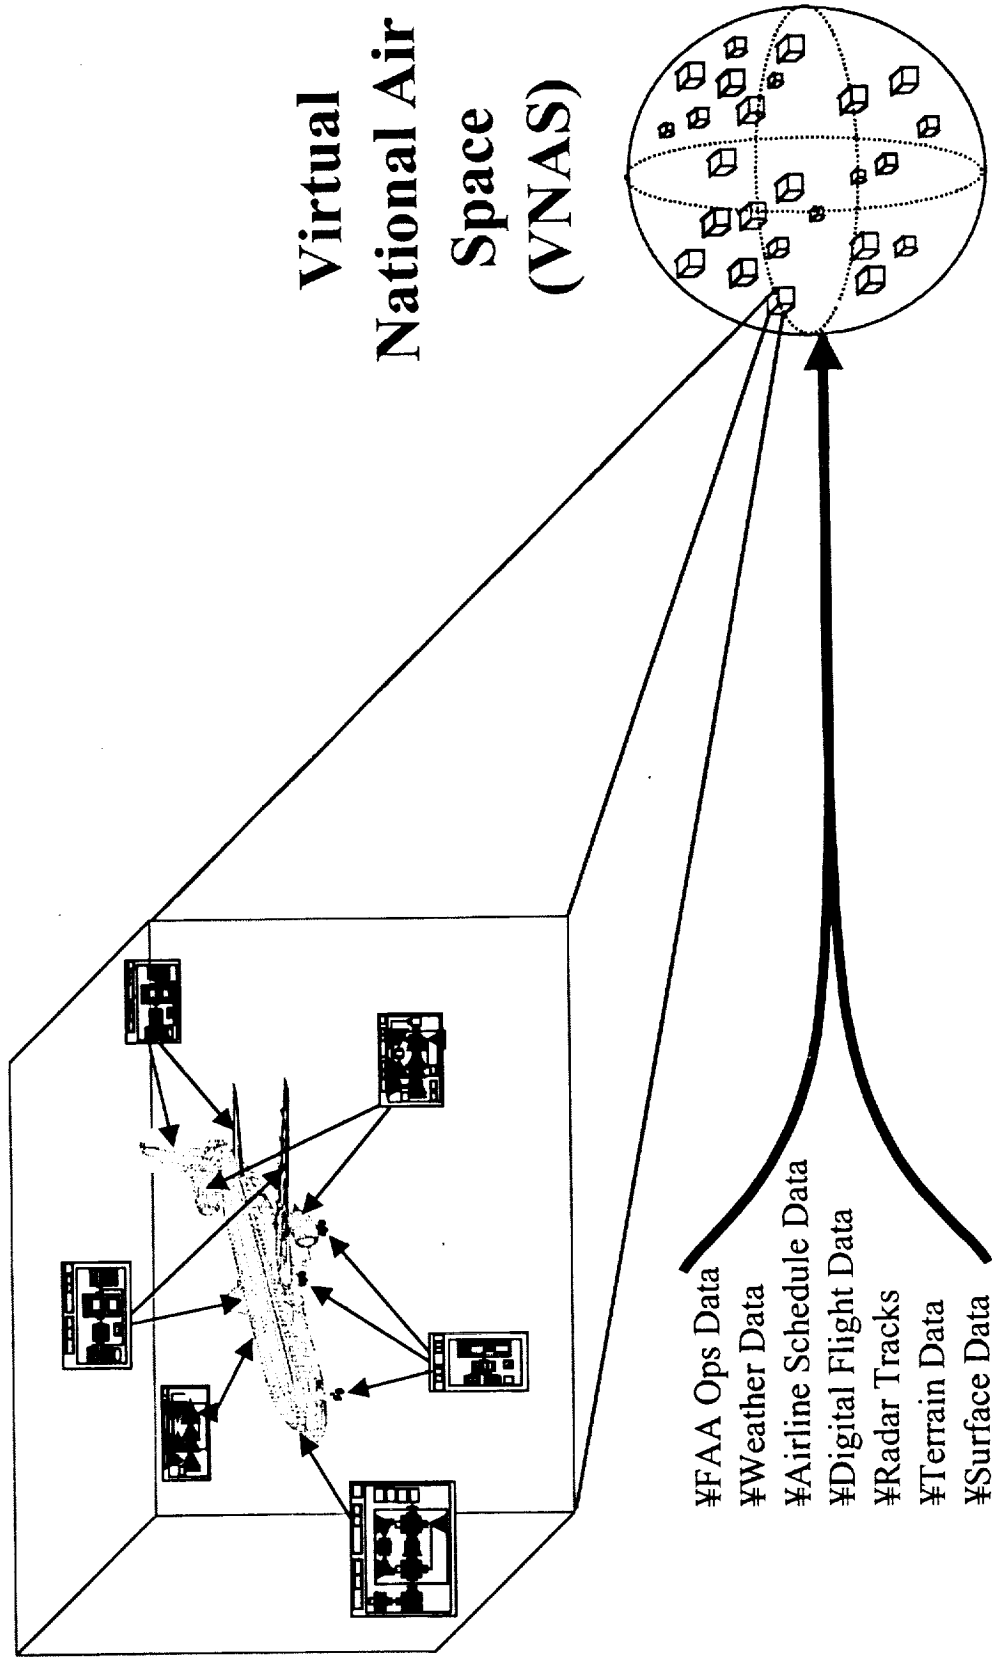
Landing Gear Models



*Whole system simulations are produced by coupling
all of the sub-system simulations*



Multi-disciplinary Simulations: Aviation Safety



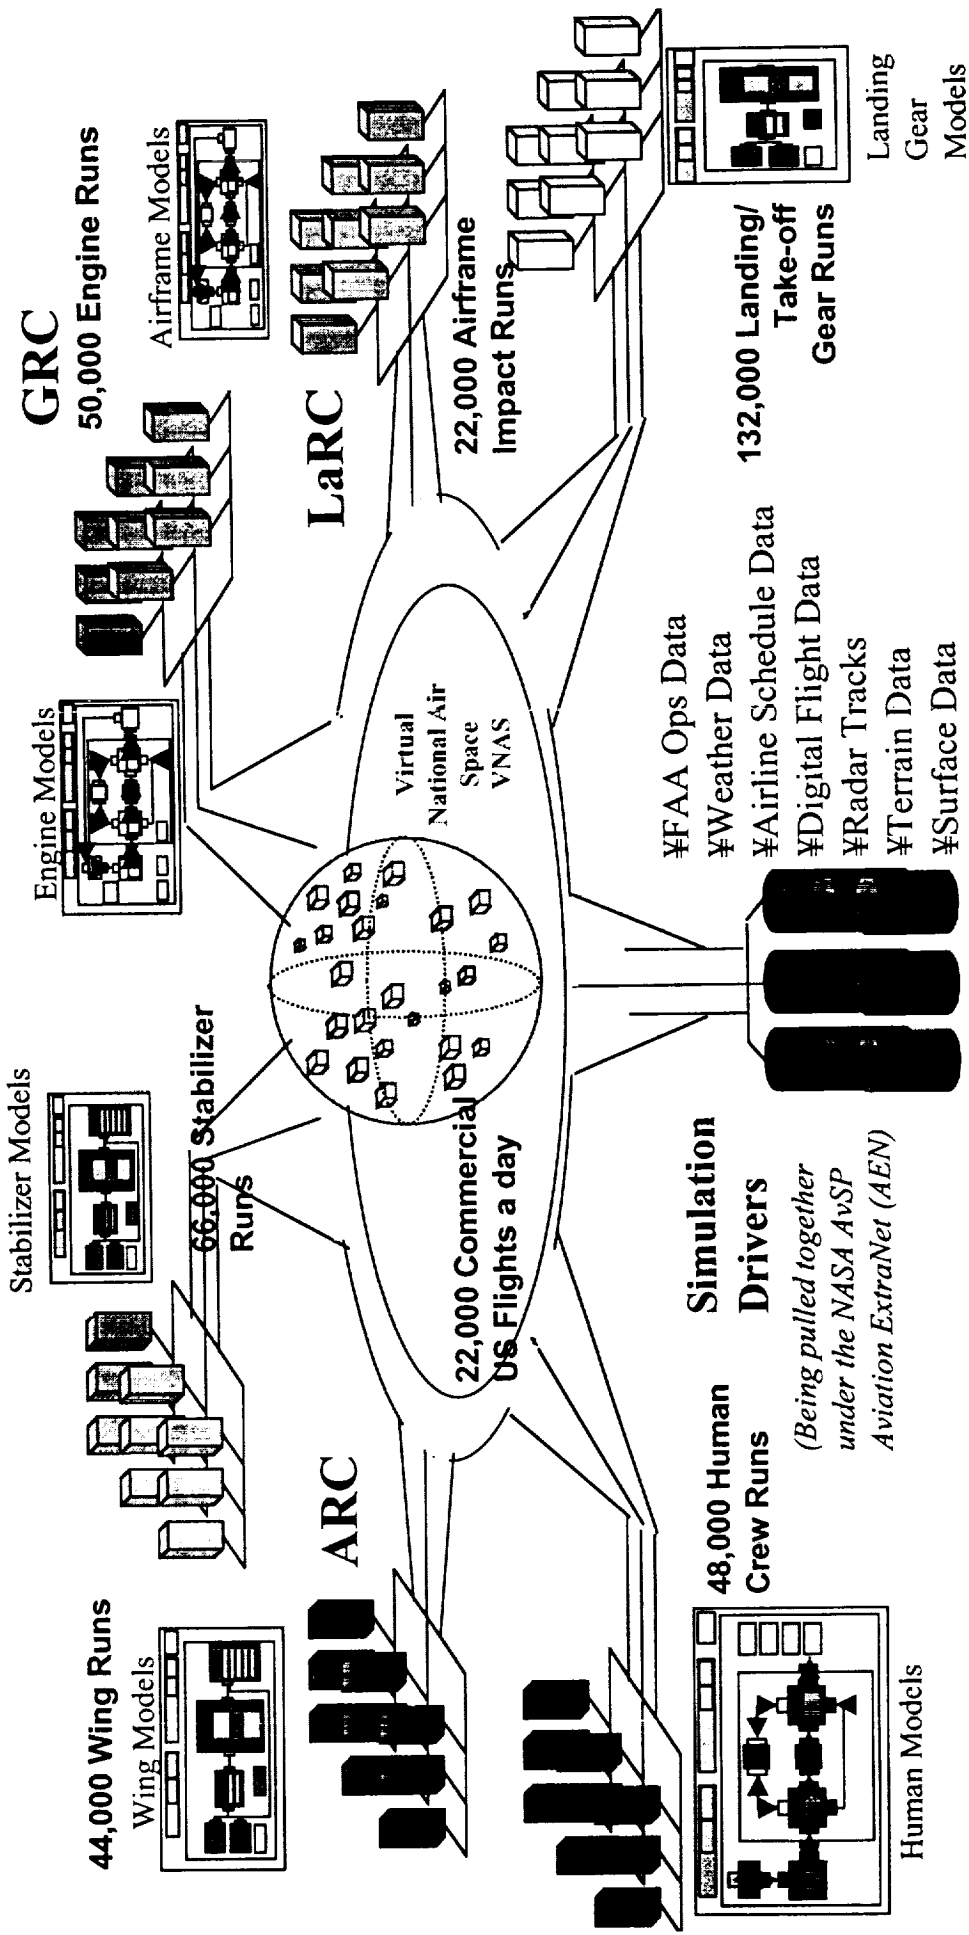
Simulated aircraft are inserted into a realistic environment



Multi-disciplinary Simulations: Aviation Safety



National Air Space Simulation Environment



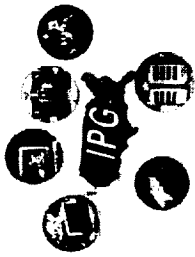
Many aircraft, flight paths, airport operations, and the environment are combined to get a virtual national airspace



Technology for Aviation Safety Simulations



- ¥ Clearly such complex, multi-component simulations will need to use aggregated computing, data, instrument, and intellectual resources across multiple NASA Centers.
- ¥ Issues for building such simulations include
 - wrapping the simulation codes so that they can be composed (e.g. CORBA and NPSS)
 - locating and coordinating resources for executing the multiple components and managing the resulting data (which is likely to be distributed)
- ¥ Grids - e.g. IPG - provide for discovery, scheduling, access, application environment construction, management, and policy enforcement for compute, data, and collaboration resource sharing and management across multiple, widely distributed, organizations





➤ ***What are Grids?***

- ¥ middleware and infrastructure for uniform, secure, and highly capable access to large-scale computing, data, and instrument resources across organizations
- ¥ tools for the dynamic construction of complex distributed applications (e.g. Virtual National Airspace)
- ¥ middleware for standardized access to data archives and standardized publication of data catalogues
- ¥ services for co-scheduling many resources to support, e.g., transient and complex, science and engineering experiments that require combinations of instruments, compute systems, data archives, and network bandwidth at multiple locations

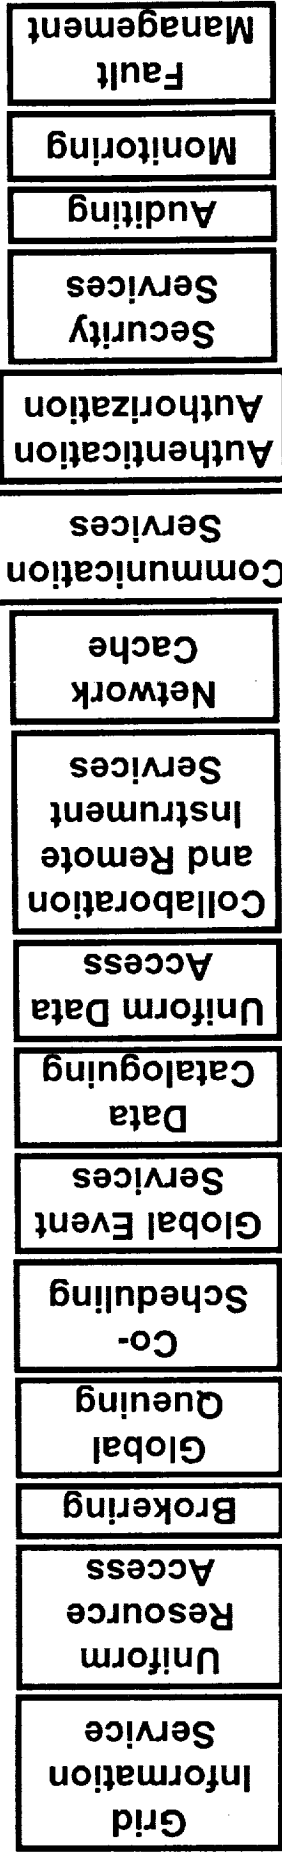


Architecture of a Grid

Discipline Specific Portals and
Scientific Workflow Management Systems

Toolkits: Visualization, data publish/subscribe, etc.
Applications: Simulations, Data Analysis, etc.

Grid Common Services: Standardized Services and Resources Interfaces



= Globus services

clusters

national supercomputer
facilities

Resources

Condor pools

tertiary storage

national user facilities

network
caches

high-speed networks and communications services

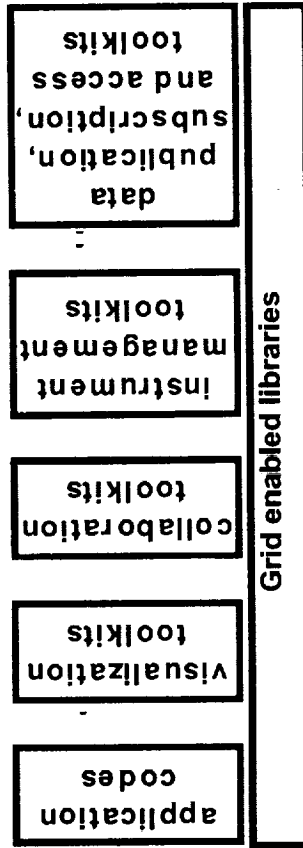


Architecture of a Grid —upper layers

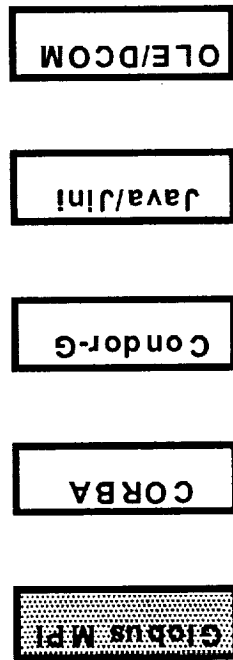
Problem Solving Environments

Tools to implement the human interfaces
Mechanisms to express, organize, and manage the workflow of a
problem solution
Access control
E.g. SciRun [24], Ecce [25], portals, WebFlow [26],...

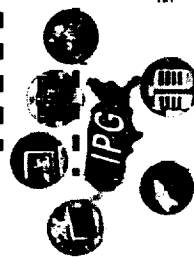
Applications and Supporting Tools



Application Development and Execution Support Services and Systems



Grid Common Services Distributed Resources



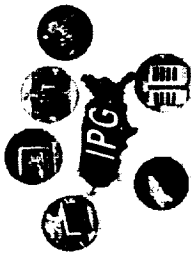


➤ *What is Being Done with Grids*



¥ Grids provide common resource access technology and operational services deployed across widely distributed virtual organizations. This allows the possibility of sharing resources, but does not automatically permit it:

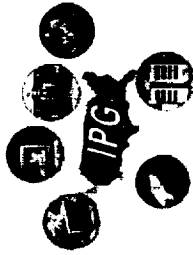
- Local authorization models are not changed by the Grid.
- Grid technology will allow common views of resources and uniform access to resources, thereby permitting very large application systems to be built, and (if policy permits) resource sharing across sites and organizations.





What is Being Done with Grids

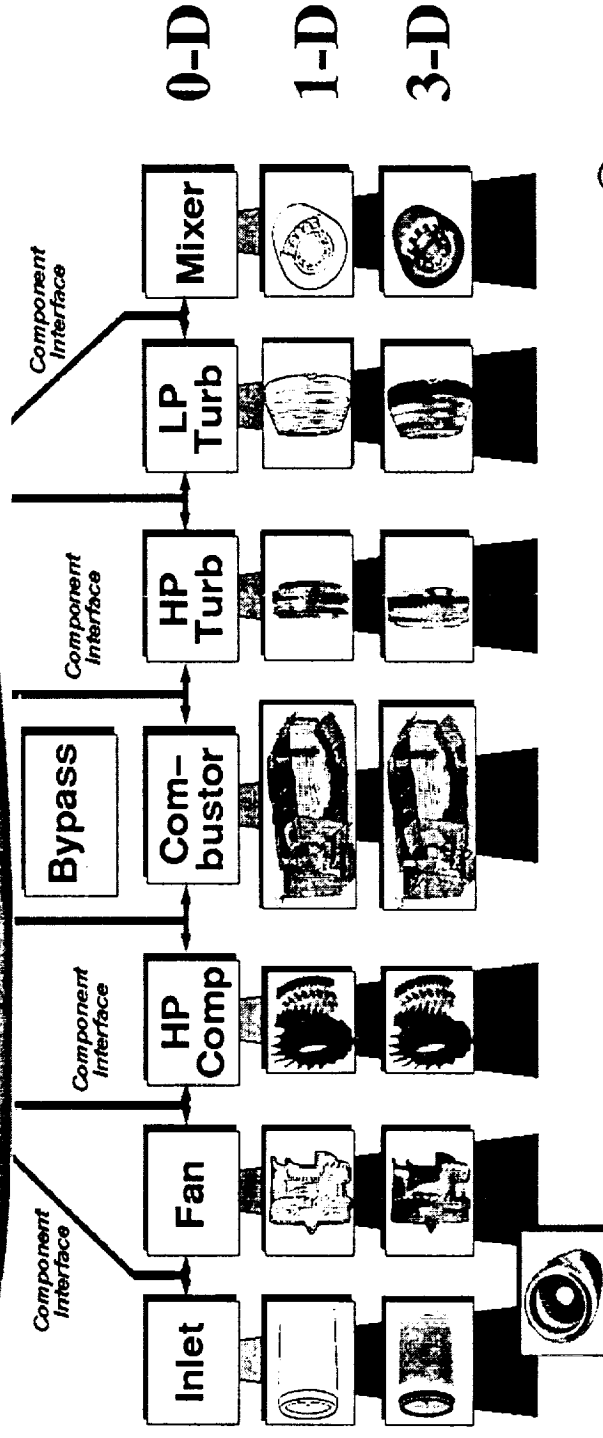
- ¥ Grids provide co-scheduling and data stream management to support large scale pipelined applications: Multi-component simulations involve executing multiple, coupled, medium to large scale simulations on multiple computing resources
- ¥ This is the basis of the Aviation Safety simulations described above



NPSS Production and Simulation Architecture



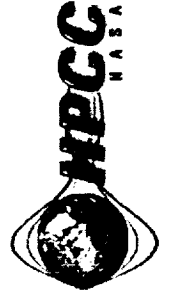
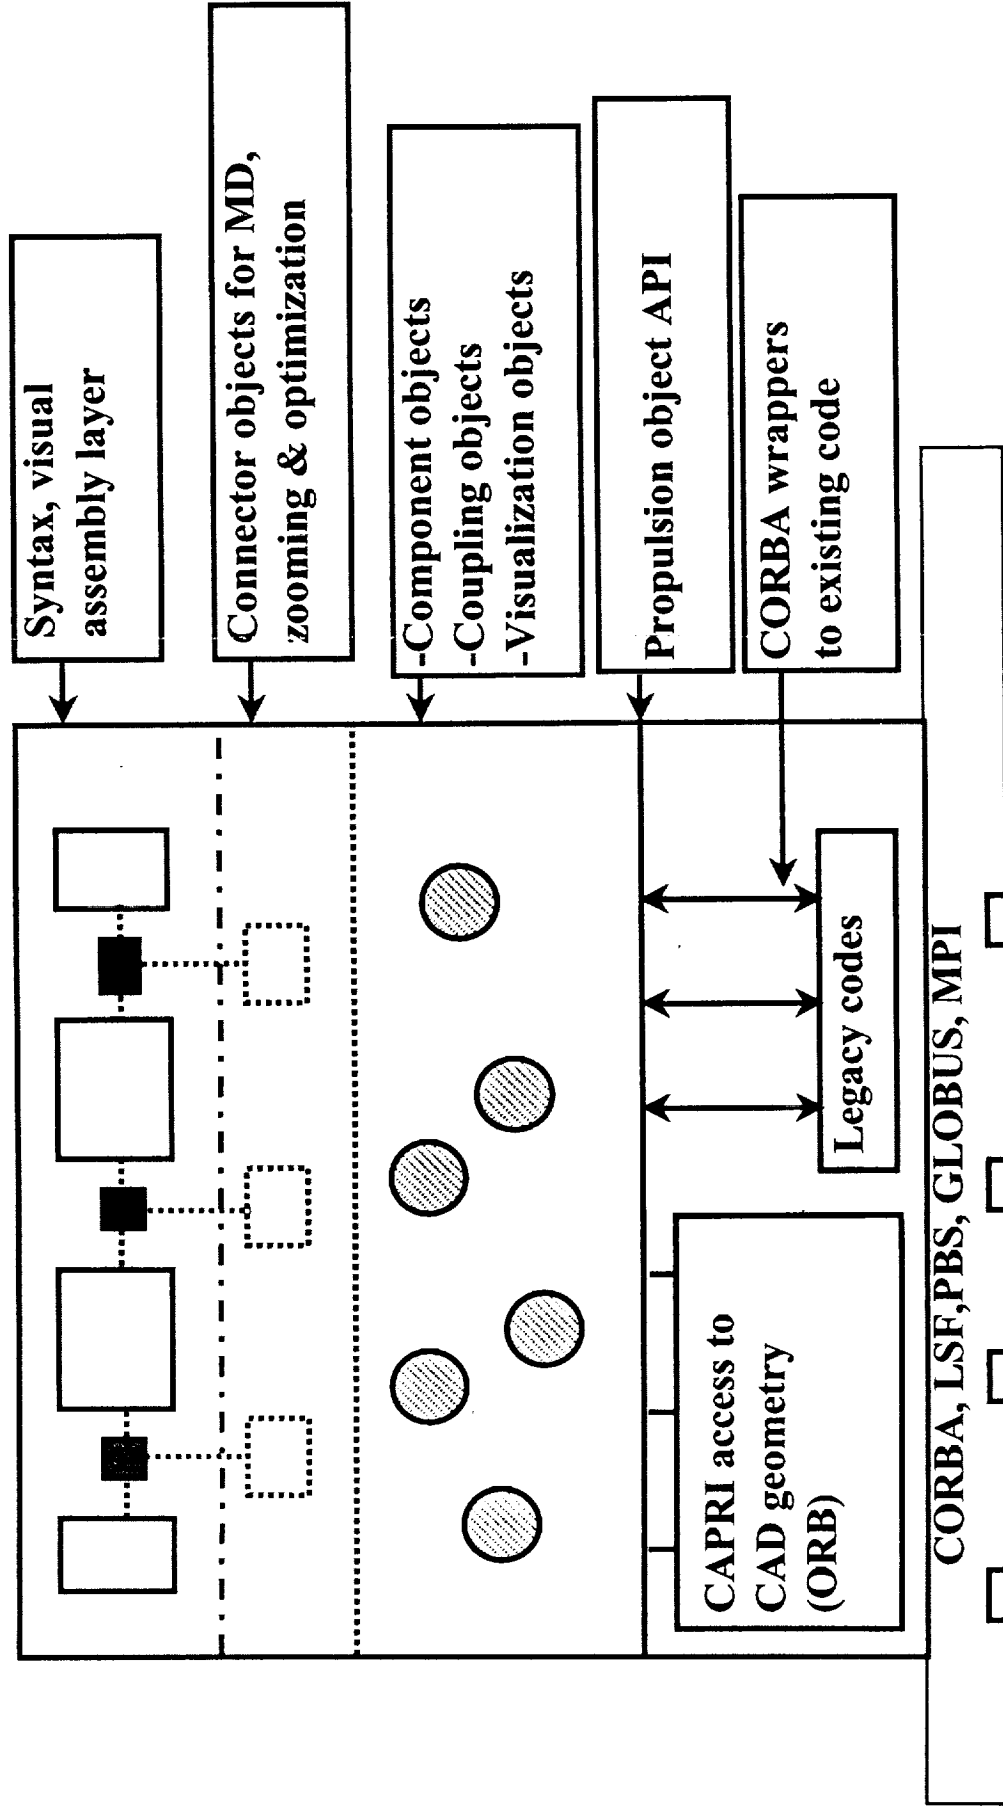
NPSS Dev. Kit
 supplies tools for
 integrating
 codes, accessing
 geometry,
 zooming,
 coupling,
 security.



NPSS slides courtesy Gregory J. Follen,
 Computing & Interdisciplinary Systems Office,
 NASA Glenn Research Center

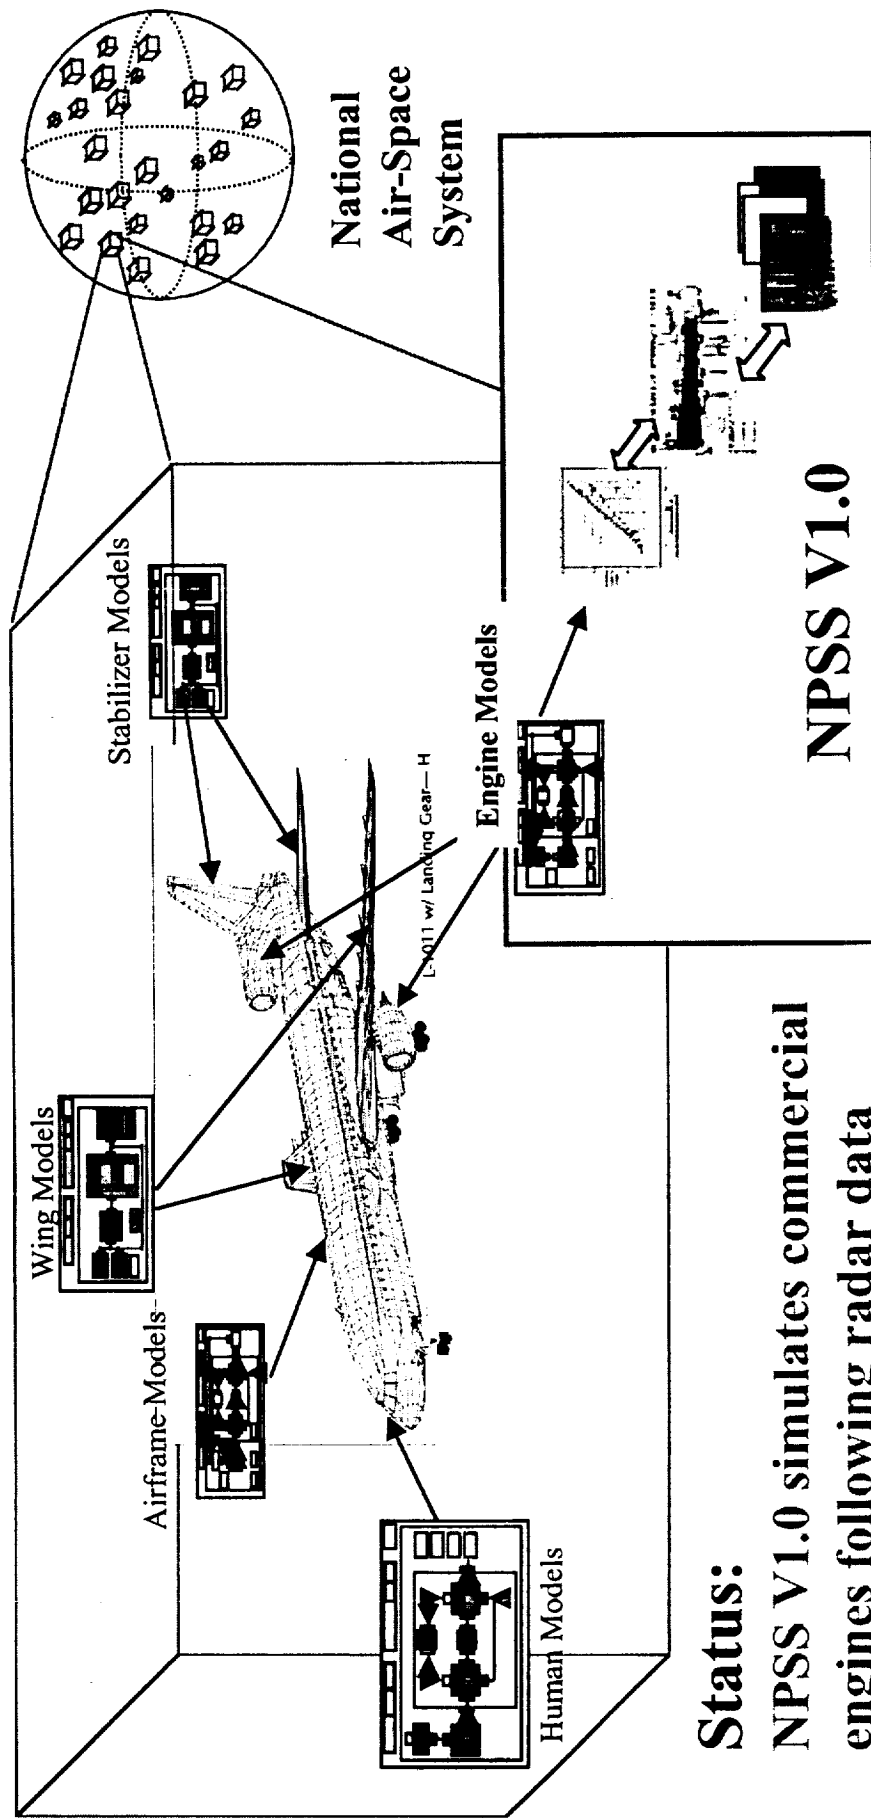


NPSS Object-Oriented Architecture



Computing and Interdisciplinary Systems Office
Glenn Research Center

Leveraging Aviation System Monitoring and Modeling



Status:
NPSS V1.0 simulates commercial engines following radar data originating from a US airport.



Computing and Interdisciplinary Systems Office
Glenn Research Center



Generic Turbofan Engine Simulation: Specify Flight and Model

Available flights (select 5 at most)

- DAL1162 (departure)
- DAL1163 (departure)
- DAL1165 (arrival)
- DAL1166 (departure)
- DAL1169 (departure)
- DAL117 (arrival)
- DAL1170 (arrival)
- DAL1170 (departure)
- DAL1174 (departure)
- DAL1175 (arrival)
- DAL1178 (departure)
- DAL1184 (departure)**
- DAL1185 (arrival)
- DAL1185 (departure)
- DAL1188 (arrival)
- DAL1189 (departure)

Selected flights:

AAL1064 (arrival)
DAL1184 (departure)

Available engine models: (select 1 model only)

- generic turbofan 0D simulation
- generic turbofan 0D simulation with 1D zooming**

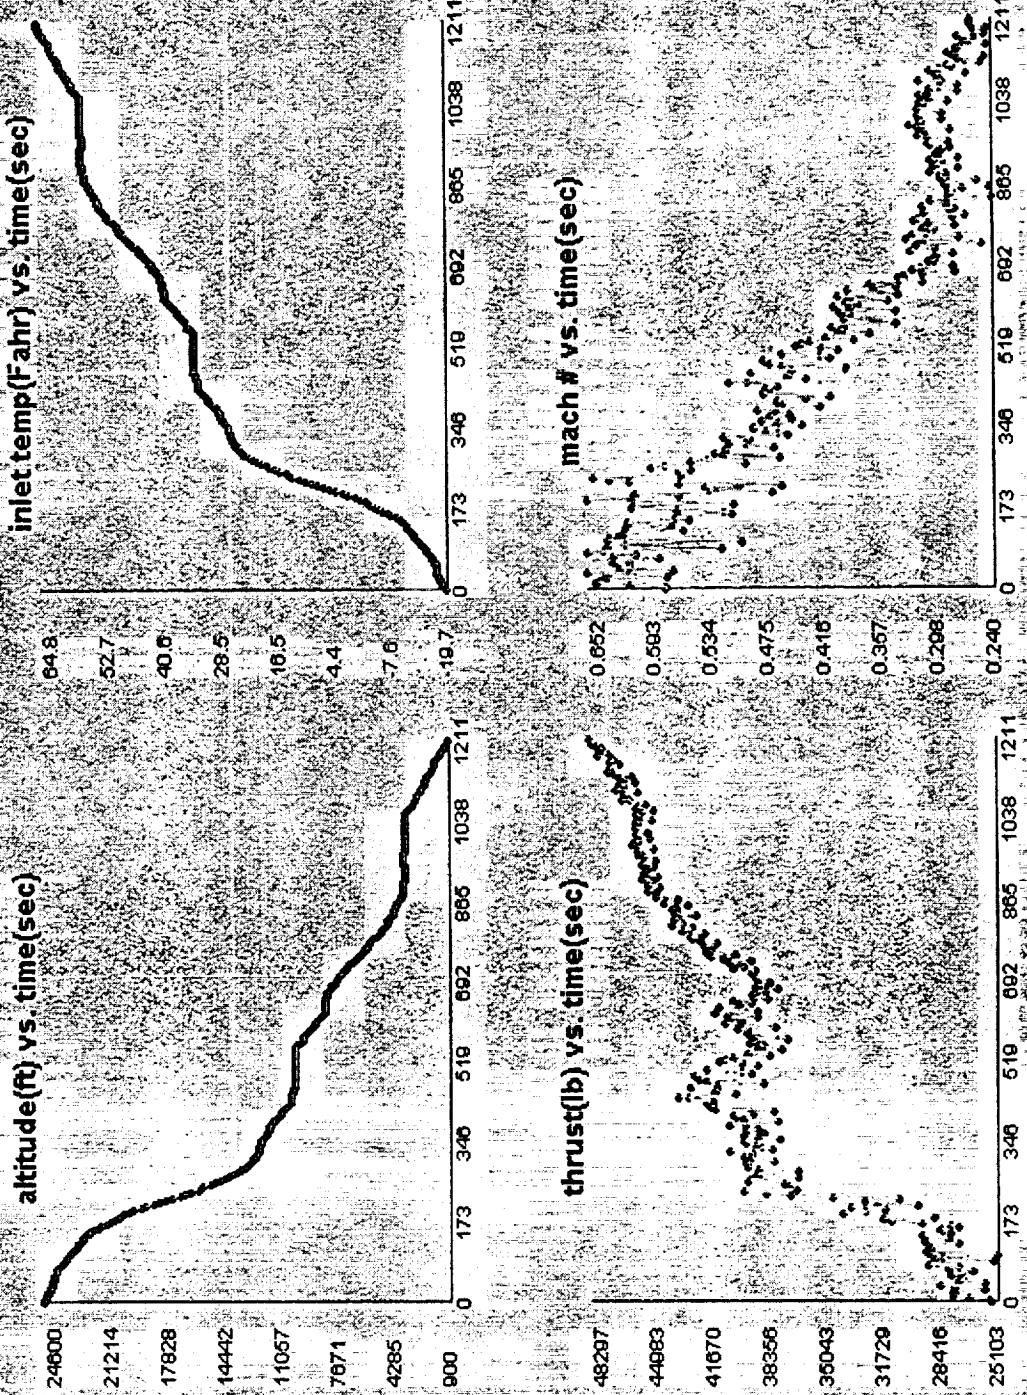


Computing and Interdisciplinary Systems Office
Glenn Research Center



NPSS Example

Generic Turbofan Engine Simulation for AAL1064 (arrival)

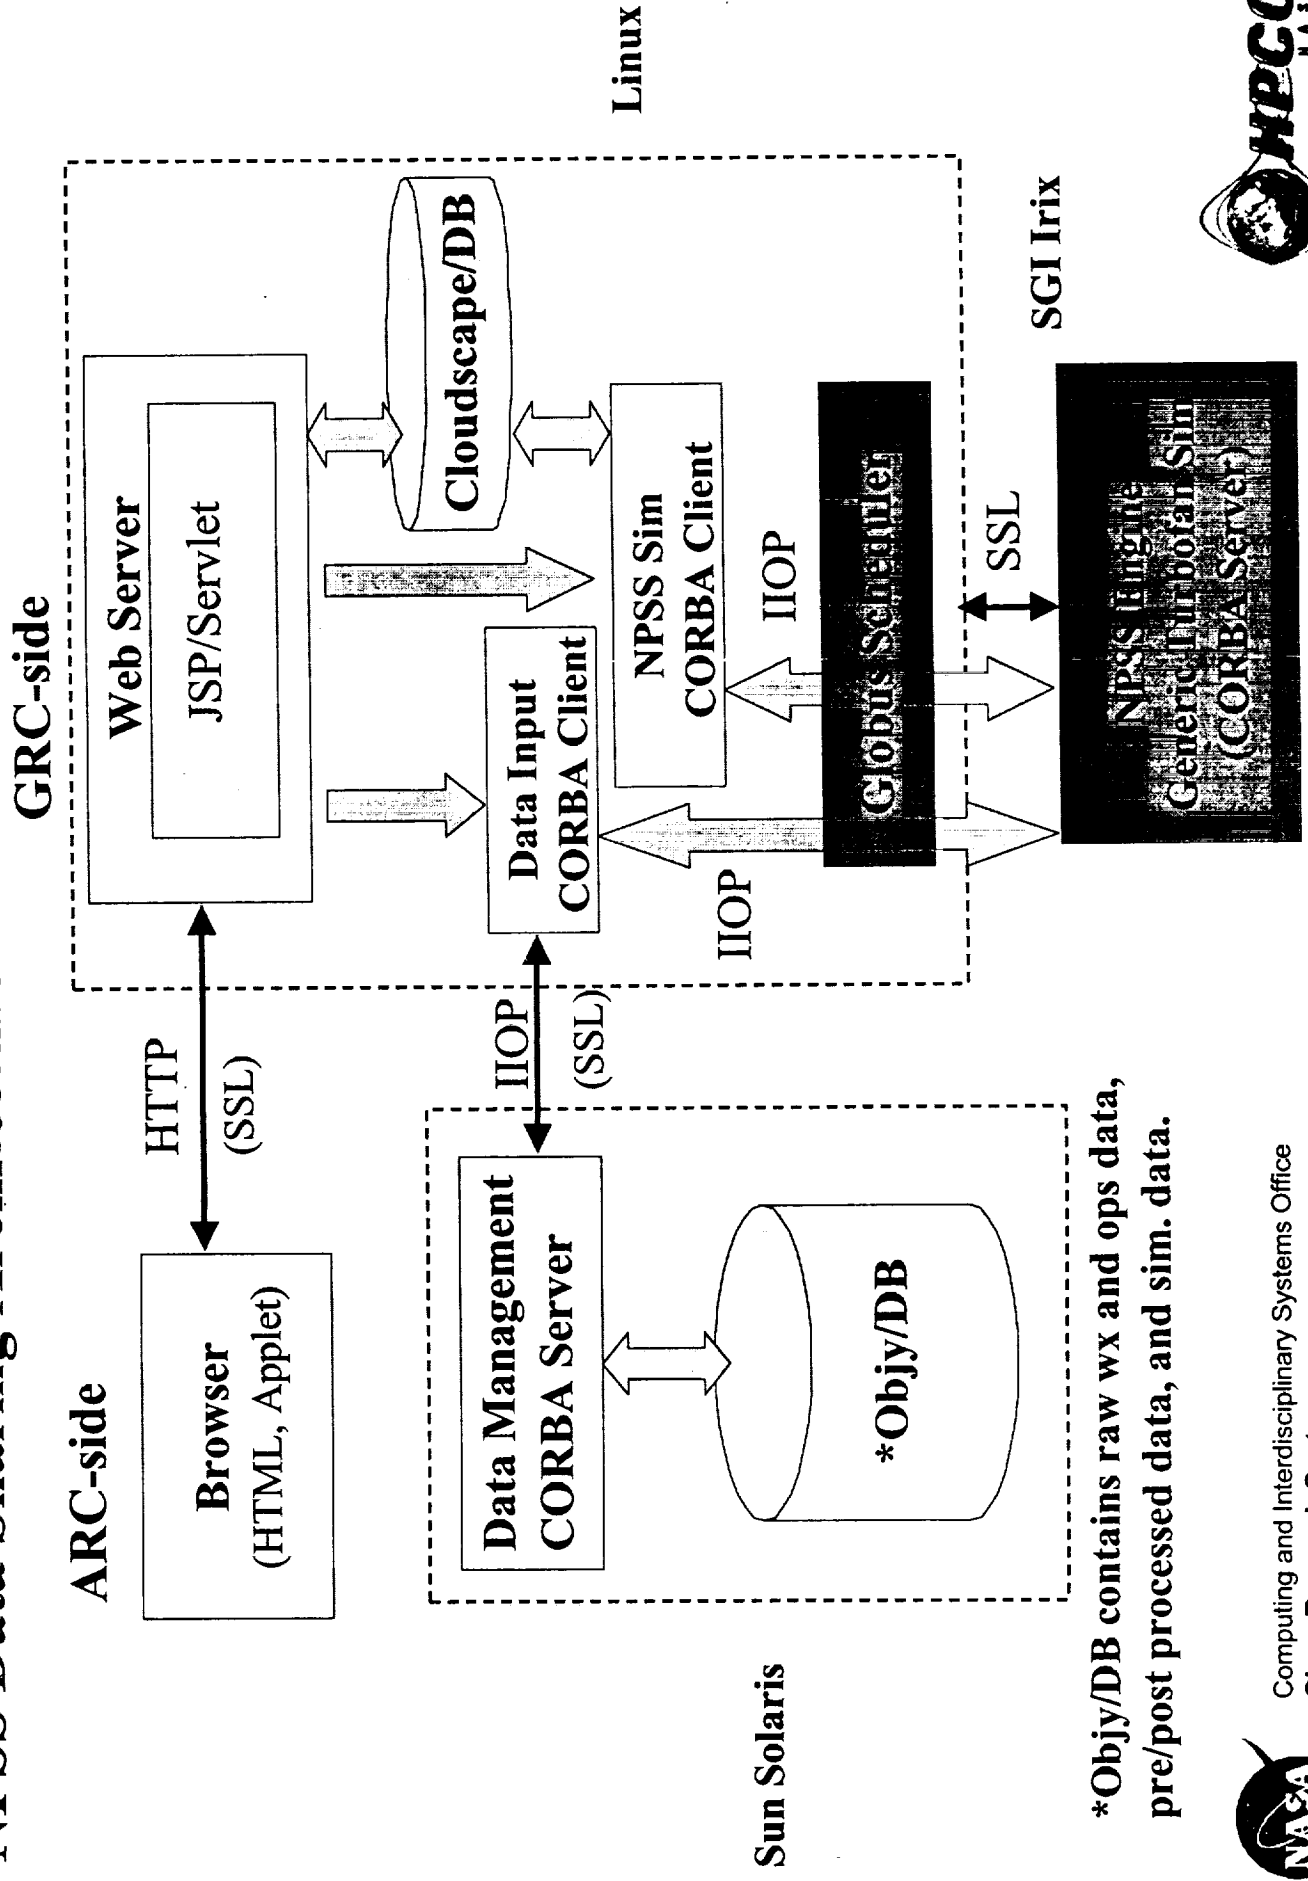


Computing and Interdisciplinary Systems Office
Glenn Research Center



NPSS Example

NPSS Data Sharing Architecture

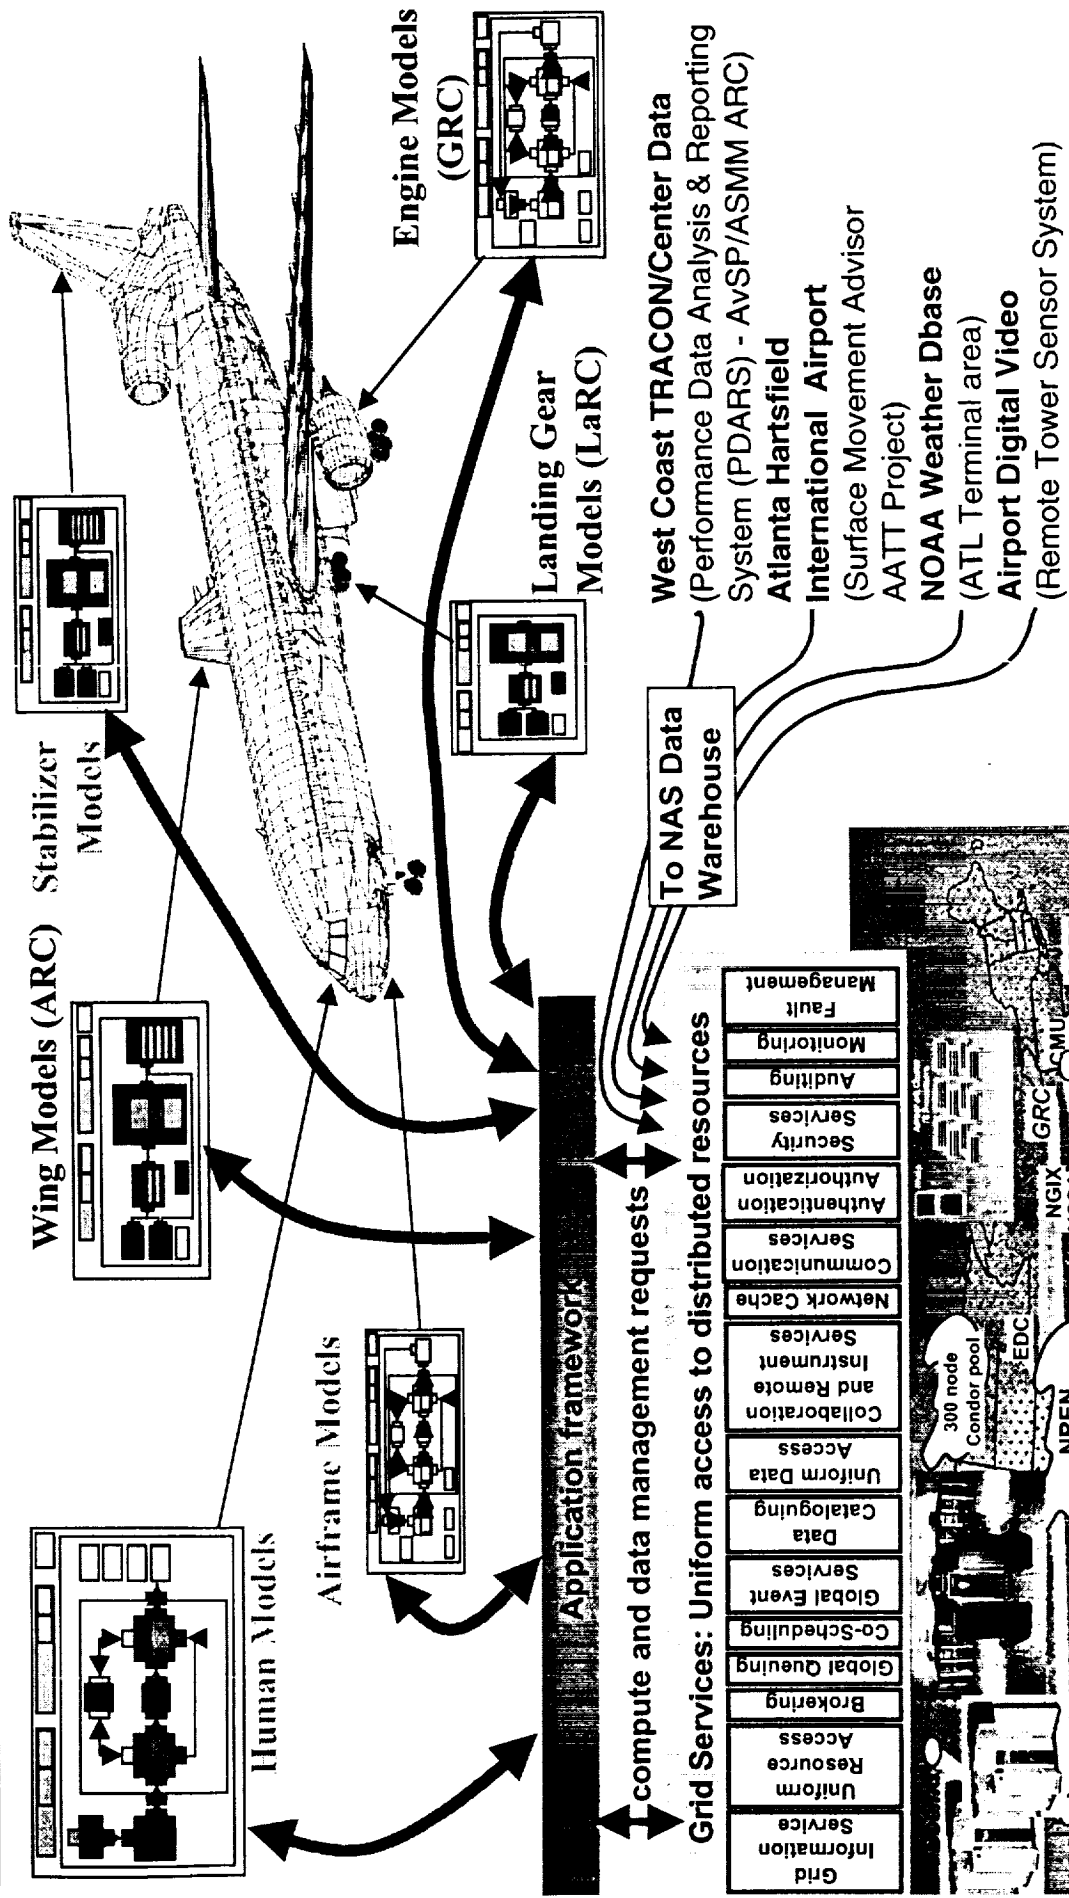


Computing and Interdisciplinary Systems Office
Glenn Research Center





Multi-disciplinary Simulations with IPG: Aviation Safety

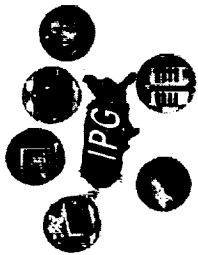


Information Power Grid managed compute and data management resources



What is Being Done with Grids (2)

¥ Grids provide the tools and middleware for discovering and access data archives that are maintained by discipline experts at many different organizations.





High Speed Distributed Data Access: IPG Milestone Completed 3/2000



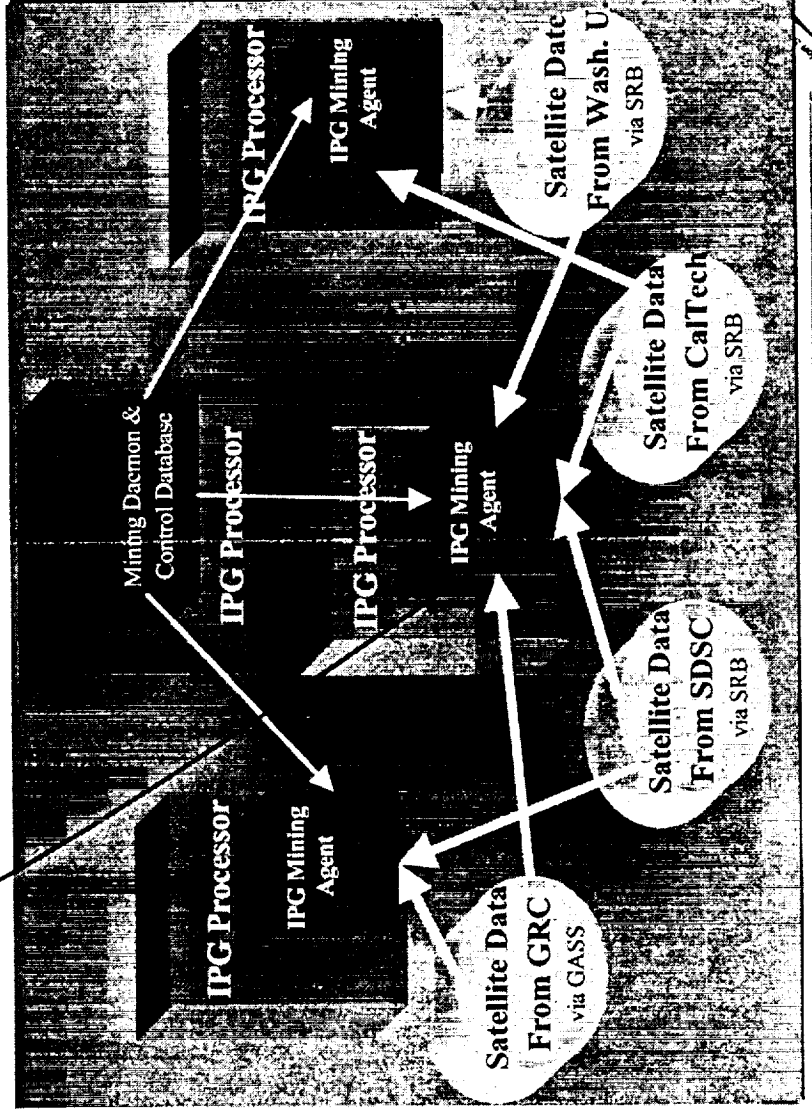
¥ Data access capabilities of IPG are demonstrated by *parallel data mining*

¥ 512 node SGI Origin at Ames uses IPG access tools to simultaneously mine hydrology data from four sites

- SDSC
- CalTech
- GRC
- Washington U.



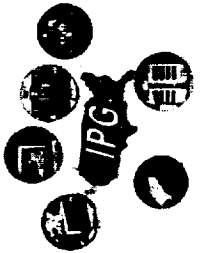
Result from one agent





What is Being Done with Grids (3)

¥ Grids provide the access and process management mechanisms for using large numbers of computing and data resources loosely coupled computations: For example, simulation parameter sweeps and certain types of experiment data analysis involve initiating and managing 100s, 1000s, and 10000s of processes.

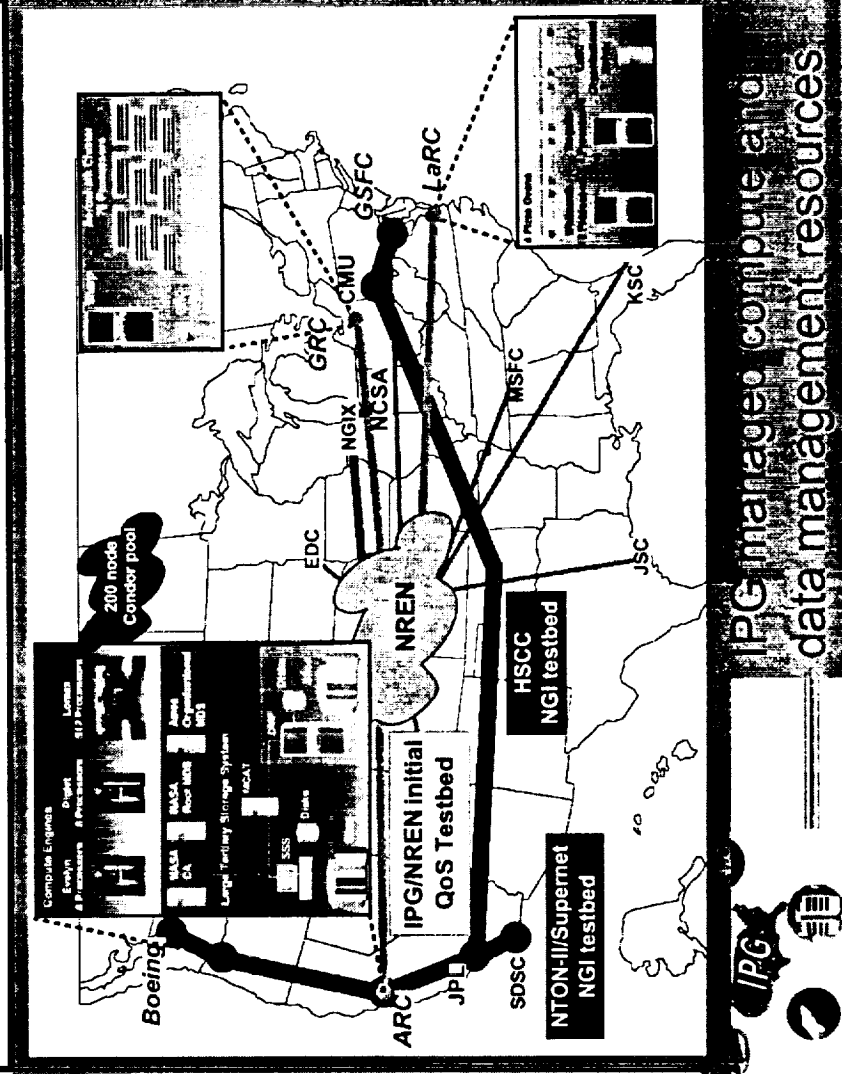
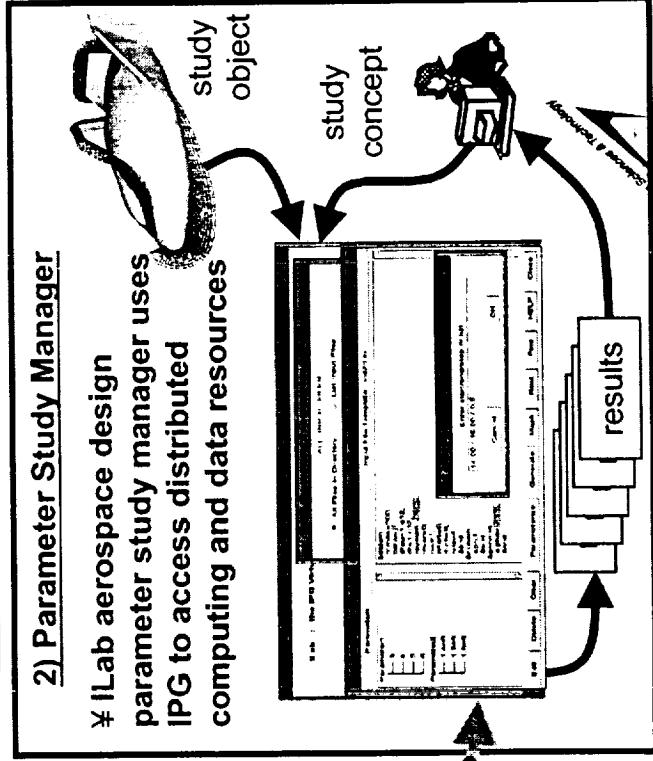
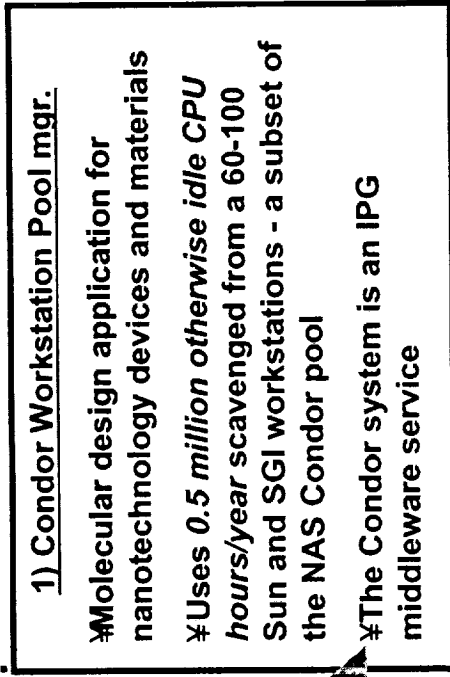
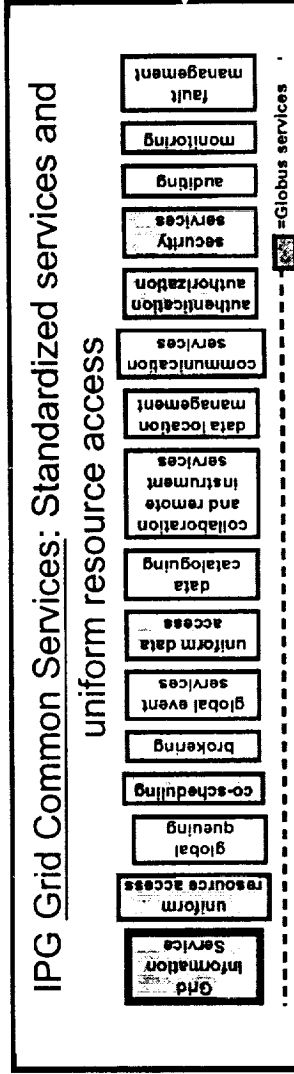




Heterogeneous Computing: IPG Milestone Completed 10/2000



- Two problem solving environments use IPG services for uniform access to heterogeneous resources.



IPG manages compute and data management resources



What is Being Done with Grids (4)

¥ In the longer term, Grids will enable certain types of very large, single problems such as overset numerical grid CFD calculations to be spread across distributed systems.

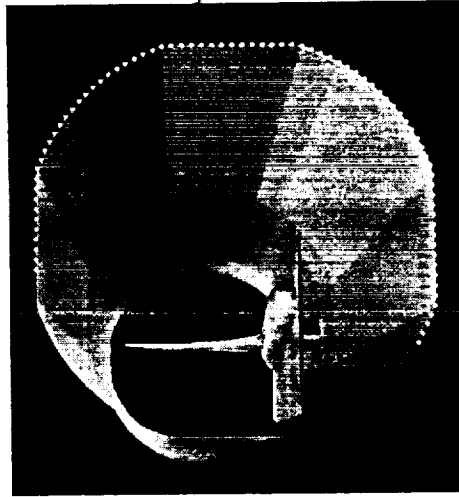
—To accomplish this we will need new approaches and algorithms that are tolerant of high and variable latency. There is R&D going on to address this issue in the long term.



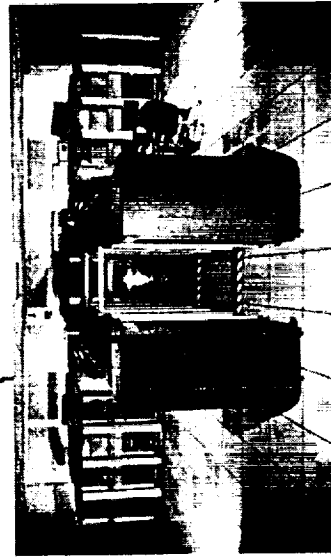


Large-scale Distributed Computing: IPG Milestone Completed 12/2000

high-lift subsonic
wind tunnel model



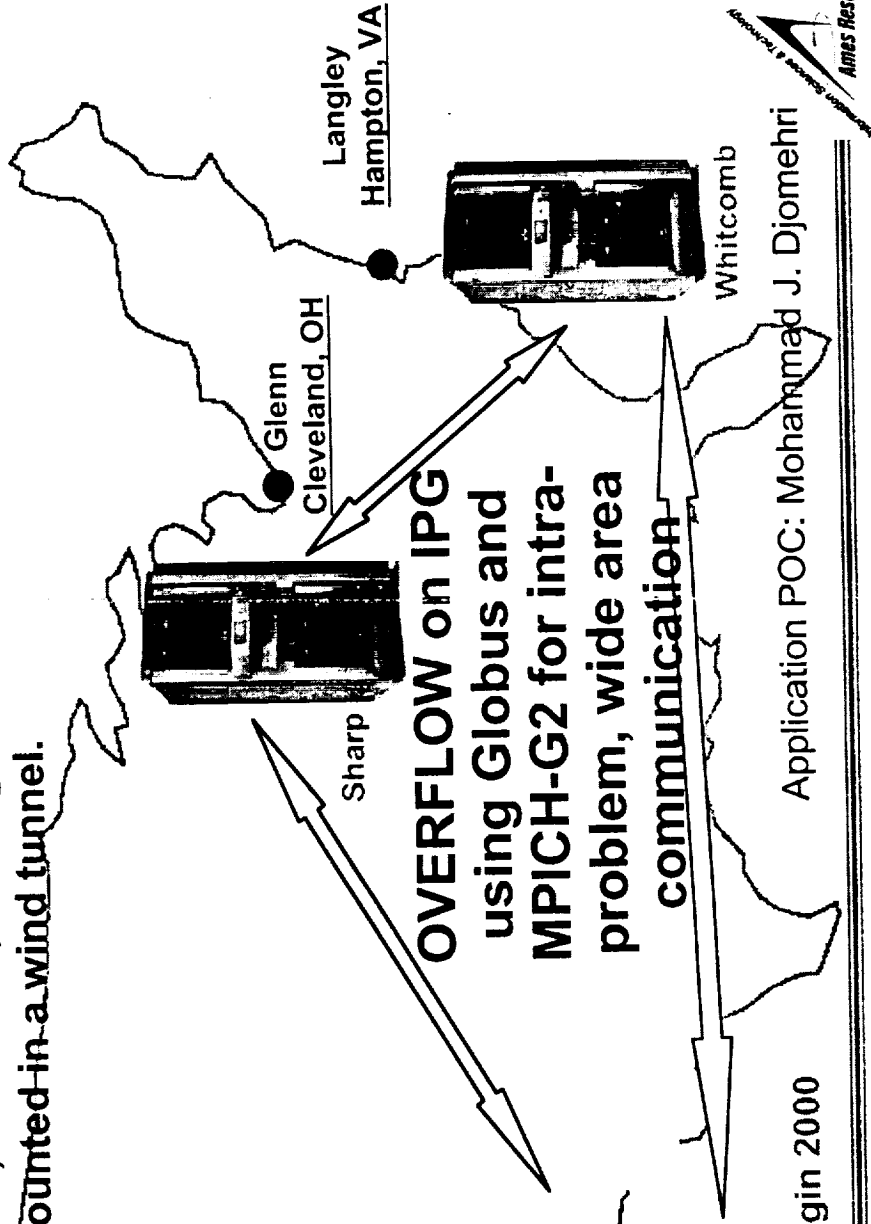
Ames
Moffett Field, CA



Lomax
512 node SGI Origin 2000



The research branch of NAS is investigating algorithms that are suitable for a Grid computing meta-platform. One candidate is overset grid codes that can tolerate timestep mis-matches on the intra-object boundaries. A version of the OVERFLOW, Navier-Stokes, CFD simulation code is being modified for this approach. It has been demonstrated operating across systems at ARC, GRC, and LaRC, solving for flow about large test objects mounted in a wind tunnel.

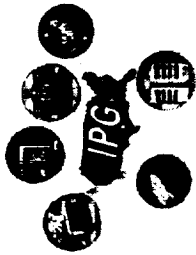




What Will Not Be Done with Grids

¥ Grids will not provide a lot of free resources.

—To produce a highly capable science Grid, organizations must place major resources —such as the 1024 node SGI Origin 2000 at Ames —on the Grid.





The Vision for IPG

- ¥ The vision for the *Information Power Grid* is to promote a revolution in how NASA addresses large-scale science and engineering problems by providing persistent infrastructure for
- highly capable computing and data management services that, on-demand, will locate and co-schedule the multi-Center resources needed to address large-scale and/or widely distributed problems
 - the ancillary services that are needed to support the workflow management frameworks that coordinate the processes of distributed science and engineering problems





➤ *Approach for NASA s IPG*



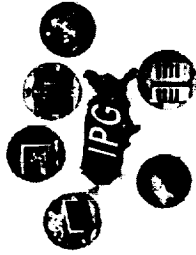
¥ Grids are built through collaborative efforts, and at the same time facilitate collaboration:

- IPG is a collaboration among several NASA Centers and the NSF Supercomputer Center consortia (PACIs), with the Grid Forum providing coordination of many institutions world wide
- The Particle Physics Data Grid is driving the development of Grid middleware to assist in the management and use of the petabytes of data generated and analyzed by the world wide high energy physics community
- the DOE Science Grid will address the software to make major instruments at the DOE Labs accessible to large-scale collaborations and allow for incorporating the DOE National Energy Scientific Computing facility into scientific experiments



Approach for NASA's IPG

¥ NASA's IPG is addressing the critical issues of how to build and operate a large-scale Grid as a production infrastructure, including providing strong security in order to address authentication, authorization, and infrastructure assurance in open science networks for both applications and Grid services





➤ *The State of IPG*



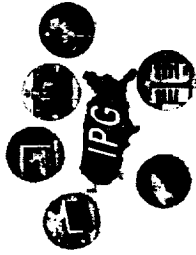
¥ Computing resources:

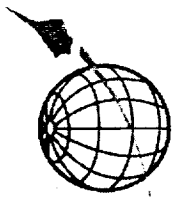
- ≈800 CPU nodes in half a dozen SGI Origin 2000s
- will add a 1024 node O2K as it completes testing and a Cray SV-1 as it completes testing (both are currently under test)
- several workstation clusters at Ames, Glenn, Langley, and JPL, with plans for incorporating Goddard
- ≈300 nodes in a Condor pool

¥ Wide area network interconnects of at least 100 Mbit/s

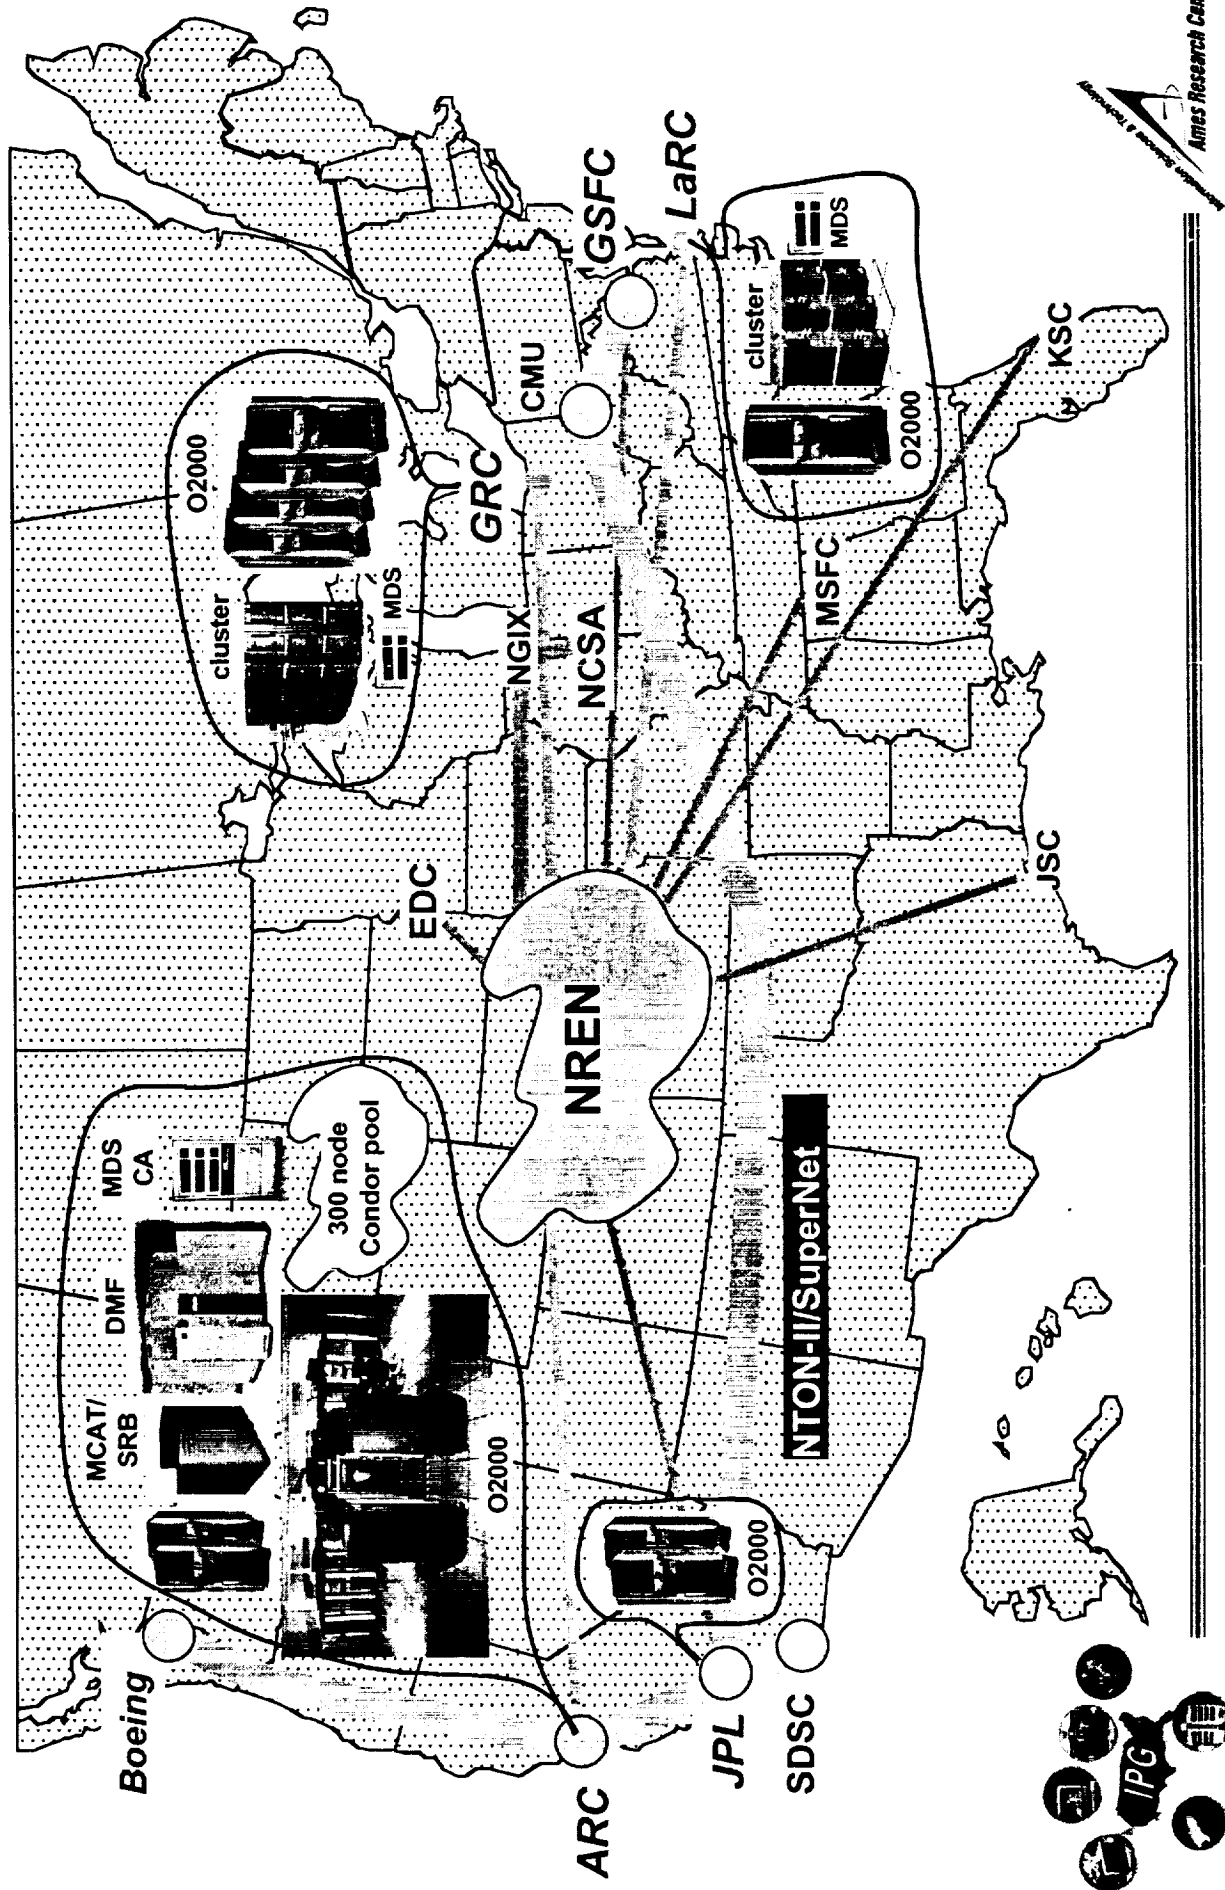
¥ Storage resources: 50-100 Terabytes of archival information/data storage uniformly and securely accessible from all IPG systems via MCAAT/SRB and GSIftp / Gridftp

¥ Globus providing the Grid Common Services





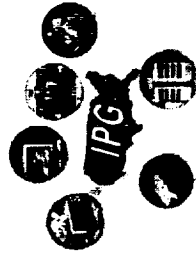
IPG Baseline System and High Data-rate Testbed





The State of IPG

- ¥ Persistent infrastructure must support Grid technologies in order to produce an operational Grid. IPG is building and operating this infrastructure, and developing and deploying Grid services for NASA.
- ¥ There are IPG operational groups at NAS whose responsibilities are:
 - Grid Information Services (the distributed master database of Grid resources)
 - Operation of the IPG computing and data systems at NAS, including the 1024 node O2K and SV-1
 - Globus software configuration and deployment
 - Grid security and authentication services
 - Grid enabled archival storage systems
 - User services
 - Condor workstation pool operation
 - The PBS batch scheduling system that provides advance reservation
 - Networking
 - Accounting

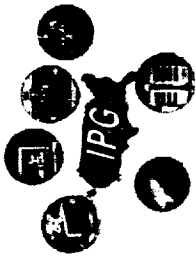




The State of IPG

¥ IPG is providing research, development, and deployment work in numerous Grid technologies:

- CORBA - Globus integration
- Integration of Legion
- CPU resource reservation
- High throughput computing
- Programming services
- Distributed debugging
- Grid enabled visualization
- Parameter study frameworks
- Network bandwidth reservation

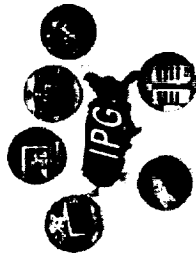




Next IPG Milestone: Online Instrumentation



¥ IPG services will facilitate real-time, or near real-time, analysis of the experiment data from on-line scientific instruments by scheduling remote high performance computers only for the times of experiment operation —this will support human steering of experiments, and/or adapting the experiment strategy while the experiment is in progress





Online Instrumentation: Real-time Experiment Interaction

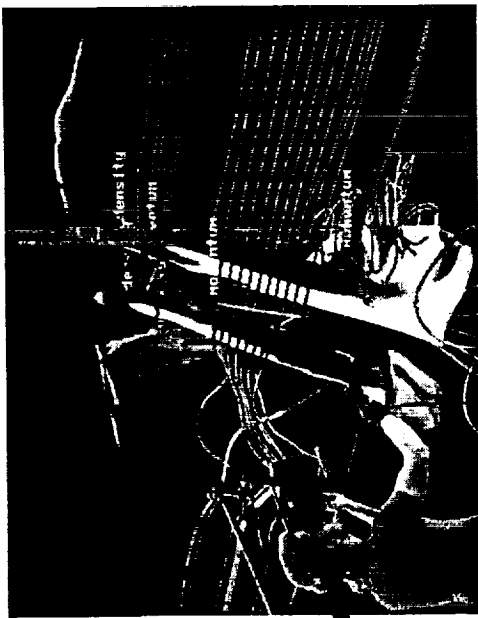
Unitary Plan Wind Tunnel



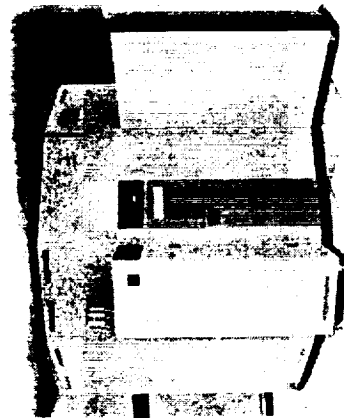
real-time
collection



multi-source
data analysis,
desktop & VR clients
with shared controls

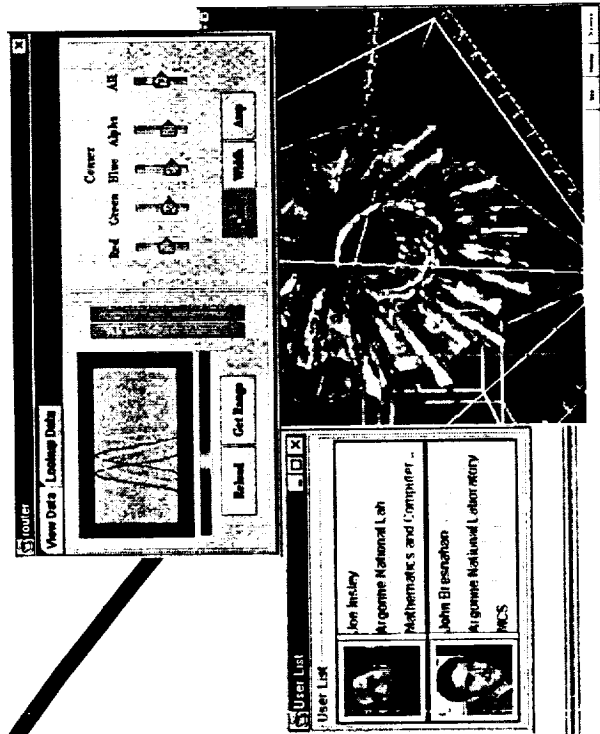
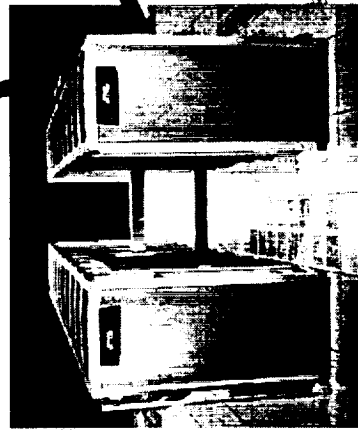


real-time experiment control



archival storage

computer
simulations



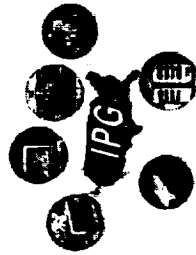


➤ *Lessons Learned for Building Large-Scale Grids*



11/10/00

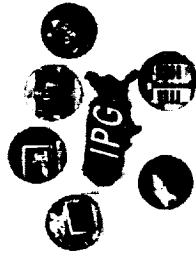
- ¥ Operational infrastructure
- ¥ Grid technology scaling issues





Steps for Building a Multi-site, Computational and Data Grid

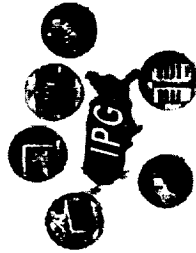
- ¥ Establish an Engineering Working Group that involves the Grid deployment teams at each site
 - schedule weekly meetings / telecons
 - involve Globus experts in these meetings
 - establish an EngWG archived email list
- ¥ Identify the computing and storage resources to be incorporated into the Grid
- ¥ Set up liaisons with the systems administrators for all systems that will be involved (computation and storage)
- ¥ Build Globus on a test system and validate the operation of the GIS/MDS at multiple sites
 - use PKI authentication and Globus or some other CA issued certificates for this test environment
 - can use OpenSSL CA to issue certs manually





Steps for Building a Multi-site, Computational and Data Grid

- ¥ Determine the model of operation for the Grid Information Service (MDS)
 - decide on Netscape LDAP hierarchy (classic model) vs. Globus OpenLDAP model
 - this may be determined by how large a Grid you plan to build
 - larger Grids may use Netscape of a meta-directory servers at the higher levels (above the GIISSs)
 - establish the GIS/resource namespace
 - ¥ be very careful about this
 - ¥ try and involve someone who has some X.500 experience
 - ¥ don't use colloquial names for institutions - consider their full organizational hierarchy when naming
 - ¥ many Grids use o=grid as the top level
 - plan for a GIS server at each distinct site with significant resources
 - get the GIS operational

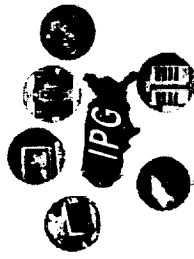




Steps for Building a Multi-site, Computational and Data Grid



- ¥ Grid Security Infrastructure (GSI) (assuming PKI based)
 - set up or identify a Certification Authority to issue Grid X.509 identity certificates
 - issue host certificates for the resources
 - count on revoking and re-issuing all of the certificates at least once before going operational
 - validate correct operation of the GSI libraries, GSI ssh, and GSI ftp
- ¥ Establish the conventions for the Globus mapfile
 - maps user Grid identities to system UIDs — this is the basic authorization mechanism for each individual platform —compute and storage
 - establish the connection between user accounts on individual platforms and requests for Globus access on those systems (initially a non-intrusive mechanism such as email to the responsible sys admins to modify the mapfile is best)

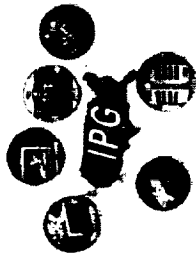




Steps for Building a Multi-site, Computational and Data Grid



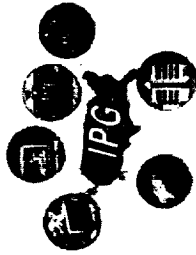
- ¥ Validate network connectivity between the sites and establish agreements on firewall issues
 - Globus can be configured to use a restricted range of ports, but it still needs ten, or so (depending on the level of usage of the resources behind the firewall), in the mid 700s
 - GIS/MDS also needs some ports open
 - CA typically uses a secure Web interface (port 443)
- ¥ Establish user help mechanisms
 - Grid user email list and / or trouble ticket system
 - Web pages with pointers to documentation
 - a Globus Quick Start Guide that is modified to be specific to your Grid, with examples that will work in your environment





Steps for Building a Multi-site, Computational and Data Grid

- At this point Globus, the GIS/MDS, and the security infrastructure should all be operational on the testbed system(s). The Globus deployment team should be familiar with the install and operation issues, and the sys admins of the target resources should be engaged.
- Next step is to build a prototype-production environment.

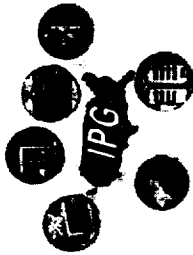




Steps for Building a Multi-site, Computational and Data Grid



- ¥ Deploy and build Globus on at least two computing platforms at two different sites. Establish the relationship between Globus job submission and the local batch schedulers (one queue, several queues, a Globus queue, etc.)
- ¥ Validate operation of this configuration





Steps for Building a Multi-site, Computational and Data Grid



¥ Establish the model for moving data between the Grid systems.

- GSIftp / GridFTP servers should be deployed on the computing platforms and on the data storage platforms
- It may be necessary to disable the Globus restriction on forwarding of user proxies by third parties in order to allow, e.g., a job submitted from platform_1@site_A to platform_1@site_B to write back to a storage systems at site A (platform_2@site_A)
- Determine if any user systems will manage user data that is to be used in Grid jobs. If so, the Grid ftp server should be installed on those systems. (So that data may be moved from user system to user job on the computing platform, and back)
- Validate that all of these data paths work correctly

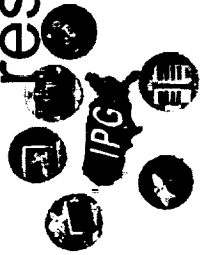




Steps for Building a Multi-site, Computational and Data Grid

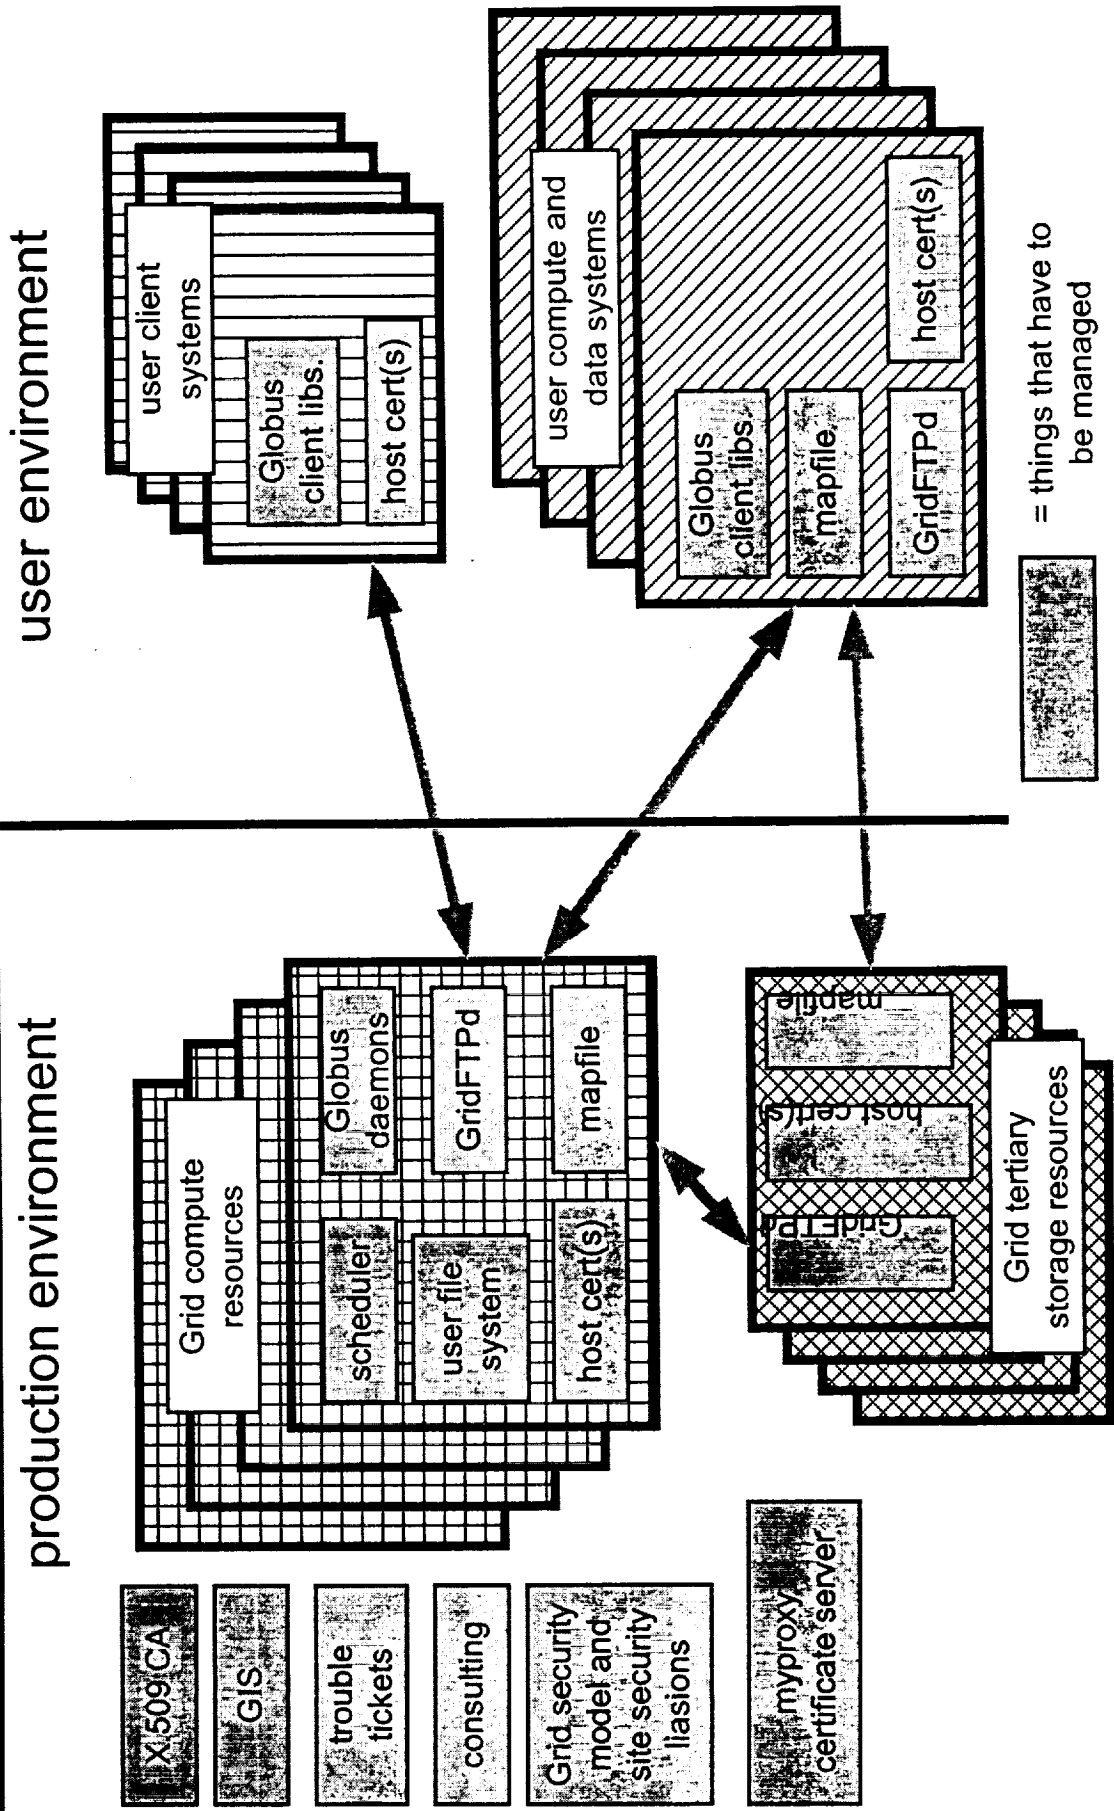


- Establish a Grid/Globus application specialist group
 - they should be running sample jobs as soon as the prototype-production system is operational
 - they should serve as the interface between users and the Globus system administrators to solve application problems
- Identify early users and have the Grid/Globus application specialists assist them in getting jobs running on the Grid
- ¥ Decide on a Grid job tracking and monitoring strategy
- ¥ Put up one of the various Web portals for Grid resource monitoring





Steps to Setting Up a Multi-Site Grid



Establish Your Grid Service Model





IPG Acknowledgements

- ¥ Almost everyone in the NAS division of the NASA Ames Research Center, numerous other people at the NASA Ames, Glenn, and Langley Research Centers, as well as many people involved with the NSF PACIs (especially Ian Foster, Lee Liming, Argonne National Lab., Randy Butler, NCSA, and Carl Kesselman, USC/ISI) have contributed to this work.
- ¥ The NASA Research and Education Network (NREN) has played a critical role in building the initial IPG.
- ¥ We would also like to specifically thank Bill Feiereisen, NAS Division Chief, and while the NASA HPCC Program Manager the initiator of IPG.
- ¥ IPG is funded primarily by NASA's Aero-Space Enterprise, Information Technology (IT) program.

